

Course	CS350 - Operating Systems
Sections	01 (11:30), 02 (16:00)
Instructor	Ashraf Abounaga
Date of Exam	February 15, 2011
Time Period	19:00-21:00
Duration of Exam	120 minutes
Number of Exam Pages (including this cover sheet)	10 pages
Exam Type	Closed Book
Additional Materials Allowed	None

Please make your answers as concise as possible. You do not need to fill the whole space provided for answers.

Question 1: (10 marks)	Question 2: (12 marks)	Question 3: (14 marks)
Question 4: (12 marks)	Question 5: (12 marks)	
Total: (60 marks)		

1. (10 total marks)

Write **T** or **F** next to each of the following statements, to indicate whether it is True or False.

- a. (1 mark) An interrupt is ignored if it occurs while another interrupt is being handled by the kernel.

- b. (1 mark) Two devices can each generate interrupts at the same time.

- c. (1 mark) Division by zero causes an interrupt.

- d. (1 mark) A thread can belong to more than one process.

- e. (1 mark) A thread executing in the kernel cannot be preempted.

- f. (1 mark) Disabling interrupts on all processors in a multi-processor system can guarantee mutual exclusion.

- g. (1 mark) A deadlock can occur with only one resource type.

- h. (1 mark) The MMU can do a TLB lookup for a running thread only if this thread is in privileged (i.e., kernel) mode.

- i. (1 mark) A software controlled TLB requires a smaller page table than a hardware controlled TLB.

- j. (1 mark) The number of pages in a program's address space is determined only by information that is in the program's ELF file.

2. (12 total marks)

a. (4 marks)

When a thread is preempted in round robin scheduling, two thread contexts are saved on that thread's kernel stack. Explain why two thread contexts are saved and not just one.

- The first context is for:

- The second context is for:

b. (4 marks)

Consider a virtual memory system that uses segmentation combined with paging, and has the following parameters:

- Each process has 8 segments.
- Each segment can be up to 8MB in size.
- Each page is 8KB in size.

Draw a diagram showing the different fields of a virtual address in this system. Make sure to label each field and indicate how many bits it contains. (Some useful information: $1K = 2^{10}$ and $1M = 2^{20}$).

c. (2 marks)

When a process makes a system call in OS/161, how does the kernel know which system call is being requested?

d. (2 marks)

In a program with multiple threads (such as the cat-mouse simulation in Assignment 1), different executions of the program with the same inputs can result in different outputs. Why?

3. (14 marks)

Consider a virtual memory system that uses paging. Virtual and physical addresses are both 32 bits long, and the page size is $4\text{KB} = 2^{12}$ bytes. A process P_1 has the following page table. Frame numbers are given in notation (recall that each hexadecimal digit represents 4 bits).

	Frame Number
0	0x00788
1	0x00249
2	0x0023f
3	0x00ace
4	0x00bcd

a. (3 marks)

For each of the following virtual addresses, indicate the physical address to which it maps. If the virtual address is not part of the address space of P_1 , write **NO TRANSLATION** instead. Use hexadecimal notation for the physical addresses.

- 0x00001a60
- 0x000051ff
- 0x00000fb5

b. (3 marks)

For each of the following physical addresses, indicate the virtual address that maps to it. If the physical address is not part of the physical memory assigned to P_1 , write **NO TRANSLATION** instead. Use hexadecimal notation for the virtual addresses.

- 0x00aceff6
- 0x00249000
- 0x00000887

c. (8 marks)

Below is the `addrspace` structure used by OS/161 when using the `dumbvm` virtual memory implementation:

```
struct addrspace {
    vaddr_t as_vbase1; /* base virtual address of code segment */
    paddr_t as_pbase1; /* base physical address of code segment */
    size_t as_npages1; /* size (in pages) of code segment */
    vaddr_t as_vbase2; /* base virtual address of data segment */
    paddr_t as_pbase2; /* base physical address of data segment */
    size_t as_npages2; /* size (in pages) of data segment */
    paddr_t as_stackbase; /* base physical address of stack */
};
```

Recall that a page in OS/161 is 4KB (2^{12}).

Below are four addresses that should not be used for `as_vbase1` or `as_vbase2`. For each of these addresses, give a reason why it should not be used.

- 0x0044f892

- 0x00000000

- 0x83410000

- 0x7fff000

4. (12 marks)

Consider two threads running in the OS/161 kernel. One thread runs the function `thread1_fn` (listed below) and the other runs the function `thread2_fn`. The functions use two semaphores `s` and `m`, which are shared between the two threads.

```
void thread1_fn (void) {
    for (;;) {
        P(s);
        P(s);
        P(m);
        kprintf("a");
        V(m);
    }
}
```

```
void thread2_fn (void) {
    for (;;) {
        P(m);
        kprintf("b");
        V(m);
        V(s);
    }
}
```

a. (2 marks)

If the initial values of the semaphores are `s = 0` and `m = 1`, list the first 10 characters of one possible output produced by the concurrent execution of these two threads.

If no output will be produced, explain why.

b. (2 marks)

If the initial values of the semaphores are `s = 2` and `m = 1`, list the first 10 characters of one possible output produced by the concurrent execution of these two threads, **different from the output that you listed in part (a)**.

If no output will be produced, explain why.

c. (2 marks)

If the initial values of the semaphores are $s = 2$ and $m = 0$, list the first 10 characters of one possible output produced by the concurrent execution of these two threads, **different from the outputs that you listed in parts (a) and (b)**.

If no output will be produced, explain why.

d. (6 marks)

For the initial values of the semaphores in parts (a), (b), and (c), are the characters that you listed the only possible characters that can be produced at the start of the output? If yes, explain why. If no, list another 10 characters that can appear as the first characters of the output, and explain what determines whether one or the other of the character sequences appear.

- part (a)

- part (b)

- part (c)

5. (12 total marks)

(Space for answer to Question 5. Question is on next page.)

List shared locks, semaphores, condition variables, or other shared global variables here (and explain the purpose of each one):

```
void thread_fn (void) {
```

(You can tear off this page of the exam if you want.)

Consider a simulation that runs in the OS/161 kernel, somewhat similar to the cat-mouse simulation of Assignment 1. In this simulation, there are N threads, and they are all of the same type. Each thread runs the function `void thread_fn(void)`. The rules of the simulation are as follows:

- There can be at most 2 concurrently running threads at any time. We will call these the *active* threads. If $N > 2$, there should be 2 active threads, and the $N - 2$ remaining threads should be waiting to become active.
- When a thread is first created, it tries to become active. If there are already two active threads, then the newly created thread waits.
- When a thread t becomes active, it calls one of two functions. If t is the only active thread, it calls the function `i_am_first()`. If there is already an active thread when t becomes active, t calls the function `i_am_second()`. Obviously, there can be no more than one thread already active when t become active, since we can have at most two active threads.
- The two functions `i_am_first()` and `i_am_second()` are part of the same critical section. These functions do not take arguments or return values, and you can assume that they are already defined.
- After the function `i_am_first()` or `i_am_second()` returns, the active thread t that called the function gives one of the waiting threads a chance to become active, and thread t waits for its next chance to become active.
- After a thread as become active 5 times (and therefore called `i_am_first()` or `i_am_second()` a total of 5 times), the function `thread_fn` exits.

On the previous page, fill in the body of the function `thread_fn` so that the above rules are followed. List the synchronization primitives and shared global variables used and explain the purpose of each one. You may use locks, condition variables, or semaphores for synchronization, but you may not use any other synchronization mechanisms (for example, you may not disable interrupts).